

26
CLAIMS

1. A method, comprising:
a master subnet manager function updating database elements of an InfiniBand
5 architecture subnet;
creating a replicated set of the database elements at each of a set of standby subnet
managers;
a standby subnet manager included in the set of standby subnet managers assuming
the master subnet manager function; and
10 the standby subnet manager included in the set of standby subnet managers that
assumes the master subnet manager function initializing the InfiniBand architecture subnet
using the replicated set of the database elements.
2. The method of claim 1, wherein initializing comprises reinitializing the
15 InfiniBand architecture subnet.
3. The method of claim 1, further comprising the standby subnet manager included
in the set of standby subnet managers that assumes the master subnet manager function
managing the InfiniBand architecture subnet using the replicated set of the database
20 elements.
4. The method of claim 1, wherein creating comprises creating the replicated set of
the database elements at each of a set of standby subnet managers out of band of the
InfiniBand architecture subnet.
25
5. The method of claim 1, wherein creating comprises creating the replicated set of
the database elements at each of a set of standby subnet managers inband on the
InfiniBand architecture subnet.
- 30 6. The method of claim 5, wherein creating inband on the InfiniBand architecture
subnet comprises creating using reliable multi-packet transaction protocol.

7. The method of claim 5, wherein creating inband on the InfiniBand architecture subnet comprises creating using reliable connection transport service.

5 8. The method of claim 5, wherein creating inband on the InfiniBand architecture subnet comprises creating using reliable datagram transport service.

9. The method of claim 1, wherein the replicated set of the database elements comprises an event subscription.

10 10. The method of claim 1, wherein the replicated set of the database elements comprises a multicast record.

11. The method of claim 1, wherein the replicated set of the database elements comprises a service record.

15

12. The method of claim 11, further comprising:
the service record comprising a lease time;
the master subnet manager function converting the lease time to a first end time;
the master subnet manager function converting the first end time to a remaining

20

time; and

the standby subnet manager included in the set of standby subnet managers converting the remaining time to a second end time, wherein the second end time is a function of the remaining time and a local time at the standby subnet manager included in the set of standby subnet managers.

25

13. The method of claim 11 further comprising:
the master subnet manager function periodically decrementing a lease time;
the lease time becoming a remaining time; and
the standby subnet manager included in the set of standby subnet managers

30

converting the remaining time to a second end time, and wherein the second end time is a function of the remaining time and a local time at the standby subnet manager included in the set of standby subnet managers.

14. The method of claim 1, wherein the replicated set of the database elements comprises an extended node record.

5 15. The method of claim 1, further comprising selecting the set of standby subnet managers based on a priority value and a globally unique identifier.

16. The method of claim 1, wherein creating the replicated set of the database elements comprises periodically updating the replicated set of the database elements.

10

17. A computer-readable medium containing computer instructions for instructing a processor to perform a method of replicating database elements in an InfiniBand architecture subnet, the instructions comprising:

15 a master subnet manager function updating database elements of the InfiniBand architecture subnet;

 creating a replicated set of the database elements at each of a set of standby subnet managers;

 a standby subnet manager included in the set of standby subnet managers assuming the master subnet manager function; and

20 the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function initializing the InfiniBand architecture subnet using the replicated set of the database elements.

18. The computer-readable medium of claim 17, wherein initializing comprises
25 reinitializing the InfiniBand architecture subnet.

19. The computer-readable medium of claim 17, further comprising the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function managing the InfiniBand architecture subnet using the replicated
30 set of the database elements.

20. The computer-readable medium of claim 17, wherein creating comprises creating the replicated set of the database elements at each of a set of standby subnet managers out of band of the InfiniBand architecture subnet.

5 21. The computer-readable medium of claim 17, wherein creating comprises creating the replicated set of the database elements at each of a set of standby subnet managers inband on the InfiniBand architecture subnet.

10 22. The computer-readable medium of claim 21, wherein creating inband on the InfiniBand architecture subnet comprises creating using reliable multi-packet transaction protocol.

15 23. The computer-readable medium of claim 21, wherein creating inband on the InfiniBand architecture subnet comprises creating using reliable connection transport service.

20 24. The computer-readable medium of claim 21, wherein creating inband on the InfiniBand architecture subnet comprises creating using reliable datagram transport service.

25 25. The computer-readable medium of claim 17, wherein the replicated set of the database elements comprises an event subscription.

26. The computer-readable medium of claim 17, wherein the replicated set of the database elements comprises a multicast record.

27. The computer-readable medium of claim 17, wherein the replicated set of the database elements comprises a service record.

30 28. The computer-readable medium of claim 27, further comprising:
the service record comprising a lease time;
the master subnet manager function converting the lease time to a first end time;

the master subnet manager function converting the first end time to a remaining time; and

the standby subnet manager included in the set of standby subnet managers converting the remaining time to a second end time, wherein the second end time is a function of the remaining time and a local time at the standby subnet manager included in the set of standby subnet managers.

29. The computer-readable medium of claim 27 further comprising:
the master subnet manager function periodically decrementing a lease time;
the lease time becoming a remaining time; and
the standby subnet manager included in the set of standby subnet managers converting the remaining time to a second end time, and wherein the second end time is a function of the remaining time and a local time at the standby subnet manager included in the set of standby subnet managers.

15

30. The computer-readable medium of claim 17, wherein the replicated set of the database elements comprises an extended node record.